FUTURE AVIONICS for X2K

BY

Dan Karmon, P.h.D
Savio Chau, P.h.D

JET PROPULSION LABORATORY
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA

ABSTRACT

Deep Space System Technology
DSST

The JPL **X2000** program is for the delivery of advanced avionics systems, cutting edge technologies, and **avionics** components to spacecraft/orbiter, **micro/nano-spacecraft**, and in-situ missions. The objective is to develop and deliver avionics systems that will serve a multitude of future missions over the next 12 to 15 years.

This program is an integral part of the NASA initiative to reduce the cost of future missions through the reduction of development cost and through capitalizing on technologies that are available in the marketplace.

The avionics system includes all the necessary elements to operate a spacecraft: Power Subsystem, Command and Data Handling, and Attitude Control.

The X2000 program is divided into three deliverables: The1st delivery for missions in the 2001 to 2003 time frame, 2nd delivery for missions in the 2003 to 2005 time frame and 3rd delivery (defined as a **System On A Chip**) in the 2006-2008 time frame.

Deliverables are based on the most advanced technologies that industry has to offer. Each delivery cycle is design to produce a meaningful technological evolvement. For example, the target for 1st to 2nd delivery growth in the Data Handling system is: A 10x improvement in volume, 20x improvement in mass, 2x improvement in power, and >10x improvement in signal/data processing. The 3rd delivery is targeted as the System On A Chip (SOAC) with yet another leap in integration levels.

In the process of development and fabrication JPL will work with industrial, academic, scientific laboratories, and other NASA center partners. This is to take advantage of the most enabling technologies and processes while maintaining broad-based synergism to reduce cost.

This paper is to describe the avionics system and to show the planned progression of system growth.

BACKGROUND

As NASA is shifting its space exploration approach from big flagships (e.g. CASSINI mission to Saturn - $1.6B) to the accelerated small - "faster-cheaper-better" missions - JPL is fast shifting gears to respond.

The Deep Space System Technology (DSST) program, otherwise known as X2000 (or - X2K) is an integral part of the NASA initiative to reduce the cost of future missions through the reduction of development cost and through capitalizing on technologies that are available in the marketplace.

Figure 1 underscores the DSST program context. It is obvious that the number of deep space missions is increasing as we embark on a new era of exploration. These new missions are predicated on the "faster-cheaper- better" NASA approach. Shrinking dollars per mission deliver a clear message that we cannot afford large individual investments, that is - doing "it" one "point-design" at a time is no longer an option.

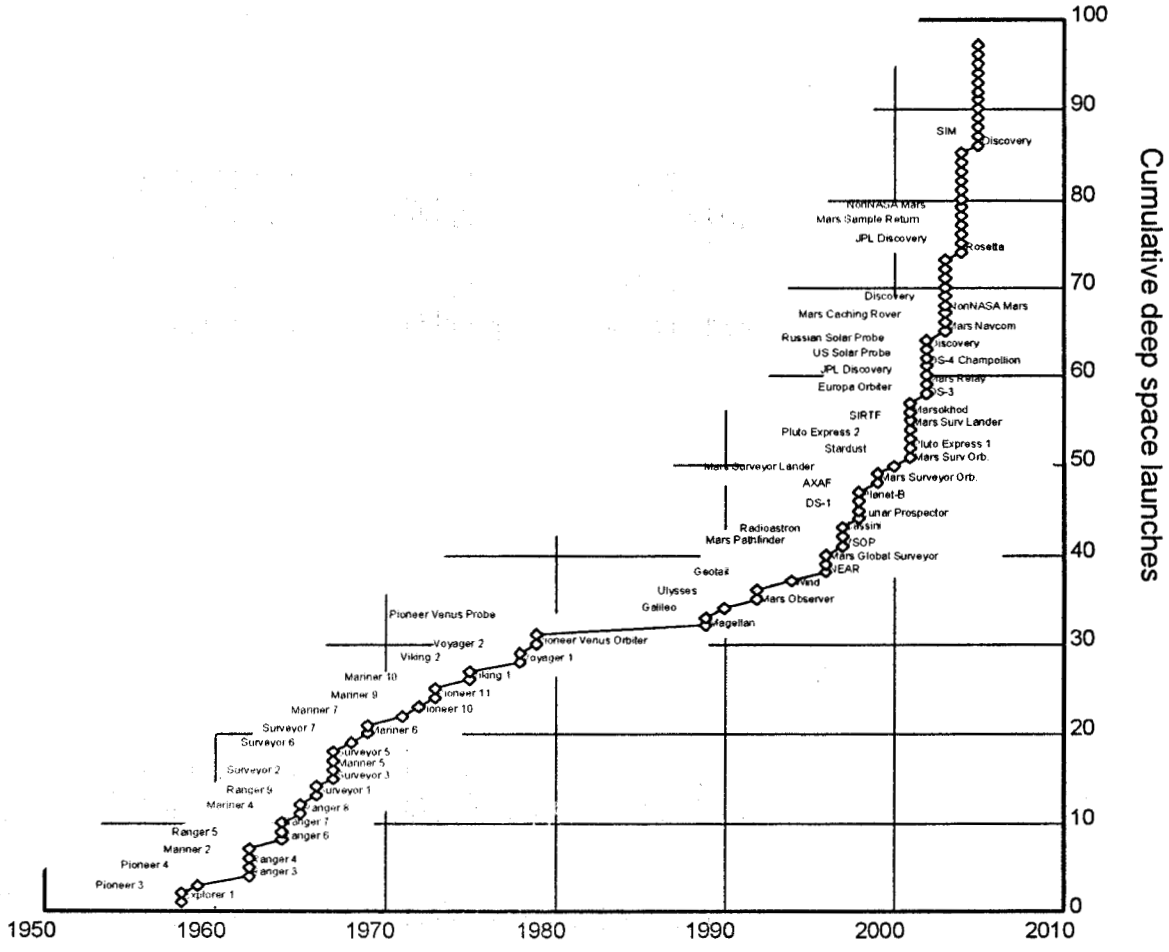

Figure 1. DSST Program Context

It is also clear that utilizing yesterday's technologies (considered safe, proven and available) in our new spacecraft will not meet the challenge, nor will the challenge be met by using old design, build, and test methods. A new way and a new approach is needed to meet the challenge of high mission density and shrinking cost.

First, let us look at the new objectives. These can be summarized as follows:

Enable deep space microspacecraft systems in the 10Kg to 50Kg.
Reduce development cost
Reduce software cost
Integrate instrument and engineering functions
Ensure a broad applicability through flexible system architecture

X2000 was conceived to meet the objective. As previously mentioned the program is currently divided into three deliverables: The1st delivery for missions in the 2001 to 2003 time frame, 2nd delivery for missions in the 2003 to 2005 time frame and 3rd delivery (defined as a System On A Chip) in the 2006-2008 time frame. This write-up to show the migration of the avionics system from one delivery to the next. All the deliverables are on an evolutionary roadmap designed to meet the objective by utilizing the best that industry can offer, by infusions of new technologies, and by "pushing the envelope".

X2000 1st DELIVERY

In FY '97, a team was formed to evaluate technologies, missions and science thrusts appropriate to the X2000. The traditional spacecraft avionics system utilized by JPL is the centralized system. It was quickly realized that such a system even when using new technologies, will not meet the flexibility expectations and will not satisfy cross cutting applicability to various missions and platforms.

It was concluded that the best approach would be to rely on distributed and scalable avionics architecture. This offers the flexibility of scalability, up and down, thus enabling spacecraft of various sizes, a variety of missions and divers science objectives.

The architecture is based on two data buses, the IEEE 1394 and the I^2C. A redundancy of two is implemented for both.

On 1394:

The IEEE 1394 bus is a commercial, high-bandwidth, standard serial bus designed for multimedia applications. The current version of the IEEE 1394 bus can support data rates of 100 Mbps, 200 Mbps, and 400 Mbps. Higher data rates will be attainable in the forthcoming IEEE 1394b specification. The most common implementation of the IEEE 1394 bus adopts a tree topology. It is a multimaster bus so that nodes have to arbitrate for the bus in order to send data. Due to the real-time nature of multimedia applications, the IEEE 1394 bus has paid special attention to guarantee that data can be delivered in time. Hence, the IEEE 1394 bus implements a technique called isochronous transaction. All the nodes requiring on-time delivery are called isochronous nodes. Each isochronous node has to arbitrate for the bus but is guaranteed a time slot (according to the allocated bus bandwidth) every 125 microseconds (referred to as an isochronous cycle), to send isochronous messages. Within each isochronous cycle, 80% of the time is available to the isochronous transactions.

While the isochronous cycles guarantee bandwidth and tightly bounded bus latency, it does not assure reliable delivery since no acknowledgment is required. Another 1394 option is the asynchronous transactions that requires acknowledgment and therefore can guarantee reliable delivery. However, the bandwidth of the asynchronous transaction is not guaranteed because it is allotted only 20% of the isochronous cycle. To avoid starving nodes, the asynchronous transaction employs a fair arbitration scheme, so that every node can send a message only once in each fair arbitration cycle. A fair arbitration cycle can span over many isochronous cycles, depending on how much of each cycle is used up by isochronous transactions and how many nodes are arbitrating for asynchronous transactions. The end of a fair arbitration cycle is signified by an arbitration-reset gap.

On I^2C:

The I^2C bus is also an industrial standard bus for many commercial applications. It is much slower then the 1394 and easier to implement. It can support data rates of 100 kbps, 400 kbps, or 3.4 Mbps. It adopts a more traditional multi-drop bus topology. The standard version of the bus can address up to 127 nodes (7-bit address) and the extended version can address up to 1023 nodes (10-bit address). The I^2C bus has two signal lines: a data line (SDA) and a clock line (SCL). Both signal lines are normally high. When a bus transaction begins, the SDA line is pulled down before the SCL line. This constitutes a start condition. Then the address bits will follow, which is followed by a read/write bit and then an acknowledgment bit. The target node can acknowledge the receipt of the data by pulling down the acknowledgment bit. After that, eight bits of data can be sent followed by another acknowledgment bit. Data can be sent repeatedly until a stop condition occurs, in which the source node signals the end of transaction by a low-to-high transition on the SDA line while holding the SCL line high.

The I^2C bus is a multi-master bus. A collision avoidance scheme is used to arbitrate conflicts among the nodes. In this scheme, all nodes can send data to the bus as long as the bus is free and no bus conflict is detected. Hence, multiple nodes can send data simultaneously as long as all data bits are the same. Once data bits from two nodes are different, the node sending a 'one' will yield to the node sending a 'zero'.

Regarding failure protection:

In addition to the standard protocol of the I^2C bus, X2000 also implemented a technique called fail-silence to detect and recover from stuck bus or babbling node. Each node in the I^2C bus has a Fail-Silent Watchdog timer which watches out for a special command called the Fail-Silent Message. A controlling bus master sends the Fail-Silent Message periodically to reset these timers. When the bus is stuck or a node babbles, the Fail-Silent Message will not be detected in time and the Fail-Silent Watchdog Timers will time out. All nodes, including the faulty node, will then disable their bus transmitters so that the bus will be free again. The receivers of the nodes will not be disabled because the nodes will have to receive bus recovery commands later on. After a wait period, the controlling bus master will start enabling each node individually. If the bus fails again when a node is enabled, then the faulty node is identified and the controlling bus master will not enable it again in the subsequent bus recovery.

As fault tolerance is for all spacecraft and space bound instrument, X2000 is also implementing a fault protection scheme. This is a five-step strategy to detect, isolate, and recover from bus failures.

Step 1: Use the fault protection features specified in the bus standards to detect, isolate and recovery from bus failure whenever it is possible. Additional fault detection mechanisms will be used to detect the bus failures that are not detectable by the standard feature.

Step 2: If a bus failure is detected but cannot be isolated or recovered by the standard fault protection mechanisms of the 1394 Bus, the I^2C Bus will be used to assist the isolation and recovery. Similarly, if bus failure is detected but cannot be isolated or recovered by the standard fault handling mechanisms of the I^2C Bus, the 1394 Bus will be used to assist the isolation and recovery.

Step 3: If both the primary 1394 and I^2C Buses failed (e.g., due to ASIC failure), the backup I^2C Bus will be activated to assist the fault isolation and recovery of the primary 1394 Bus.

Step 4: If the backup I^2C Bus failed, the backup 1394 Bus will be activated to assist the fault isolation and recovery of the primary 1394 Bus. If the primary 1394 Bus has too many failed nodes (i.e., bus partitioned by the failed nodes or the number of hops exceeds the limit), the backup 1394 Bus will take over.

Step 5: If all the above steps failed, each node will enable and disable its ports on both buses according to a distributed recovery procedure and try to communicate with its neighbors.

The X2000 team also developed the concept and detail for the various avionics sub-components for the 1st delivery. "SLICES" of various avionics components are connected to the data buses as needed. Some of these "slices" are:

SFC - System Flight Computer
PCS - Power Control Slice
PSS - Power Switch Slice
MCS - Micro Controller Slice
NVM- Non-Volatile Memory

Each component, or slice, takes full advantage of cutting edge technology and processes that are available in the marketplace. For example, the SFC takes advantage of the Power PC, 0.35 Micron SOI (Silicon on Insulator) process, and radiation hard designs. The power elements (slices) take advantage of HFET and HDI (High Density Interconnect). In addition, reusable designs are used extensively: The bus interface ASICs are comprised of commercial off-the-shelve synthesizable 1394, I^2C and PCI cores (referred to as Intellectual Properties or IP). This Approach is the underlining principle of the X2000 program, which is primarily a technology program.

A simplified system block diagram for X2000 1st Delivery is shown in Figure 2 below. This program, which is currently in development, is to deliver flight hardware to a number of flight programs (missions) in the year 2001 time frame. EUROPA and CHAMPOLLION/ST4 are the current flight mission customers with a number of additional missions negotiating for flight hardware. Figure 2 depicts the generic concept. The specific mission related details are not depicted.

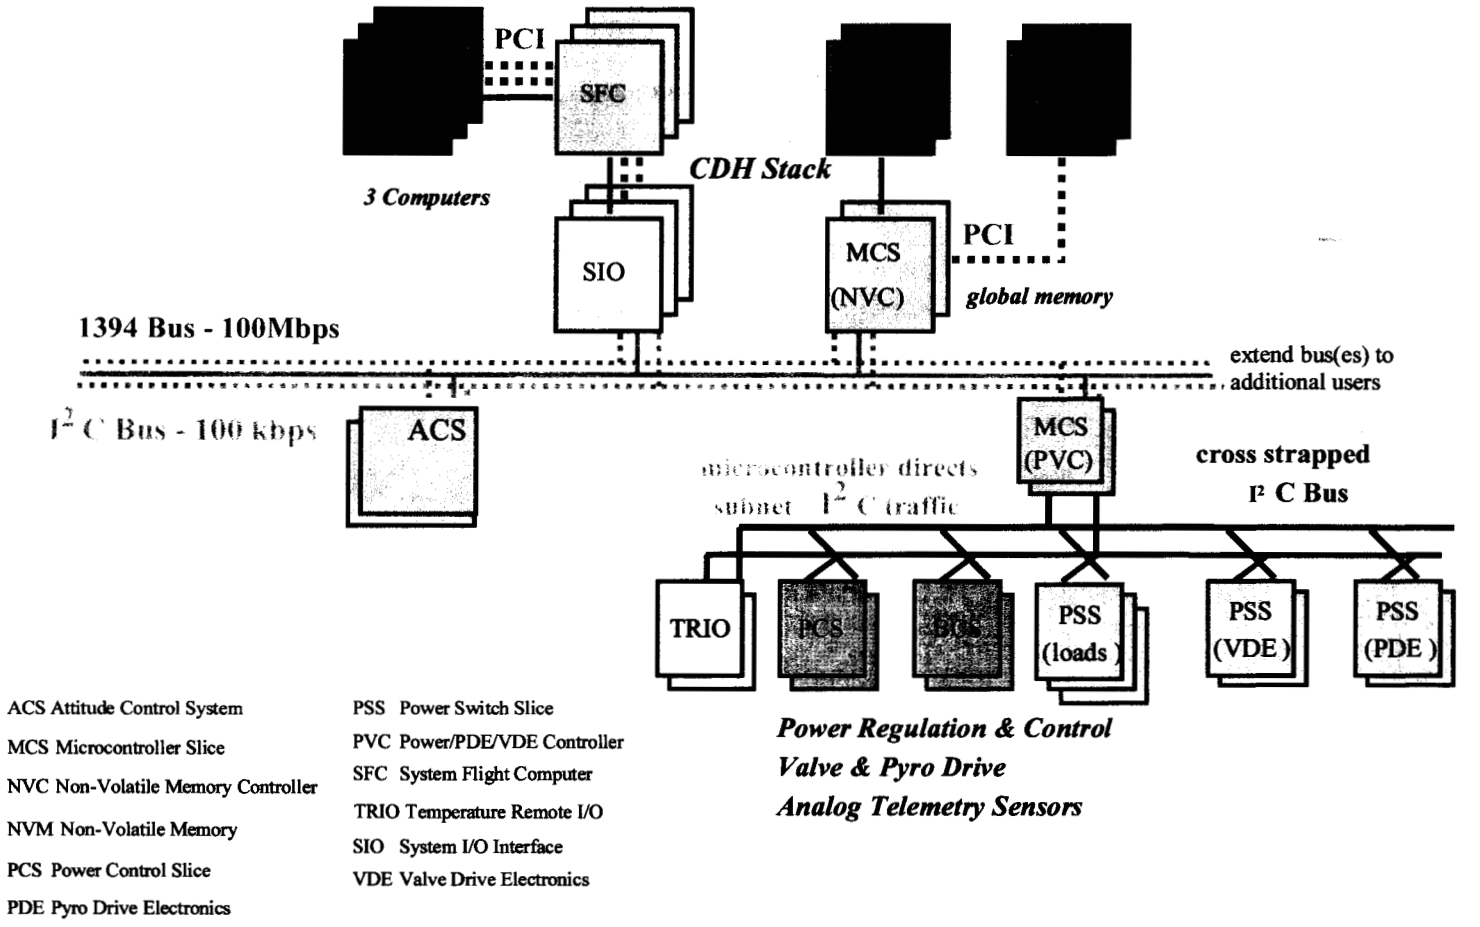

FIGURE 2. X2000 1ST DELIVERY- A SIMPLIFIED BLOCK DIAGRAM

The X2000 1st delivery avionics "slices" are designed for a Compact PCI card cage and back-plane. The program is currently in the detailed design and "contract letting" phase.

SYSTEM ON A CHIP (SOAC) TECHNOLOGY PROGRAM

As previously mentioned, X2000 is to deliver "advanced avionics systems, cutting edge technologies.....", and is currently divided into 1st, 2nd , and 3rd deliveries. One of the key X2000 supporting technology-programs resides in the JPL Center for Integrated Space Microsystems (CISM.). It is the System On A Chip (SOAC) program. SOAC is the technology pipeline for the X2000 deliveries.

SOAC is currently working on a wide range of enabling technologies in partnership with universities and industry. A list of the major thrusts is depicted below.

Micro Inertial Measurement System (micro gyro)
PMAD components
Active Pixel imagers
RF Communication front end
Ultra low power electronics
Wide temperature electronics
Embedded passives
Integrated sensors
System integration - ASIC technology
SOAC reliability

Of particular interest to avionics are sensor integration and analog and mixed-signal design. The objective is to develop system and circuit concepts that exploit potentials offered by advanced deep-submicron space-qualified fabrication technologies and to establish infrastructure for developing SOAC systems for future spacecraft applications.

Some of the related and infrastructure tasks are:

Identify, develop and demonstrate CAE tools and procedures that support the design, fabrication, rapid prototyping and test of advanced modular SOAC designs.
Establish commercial and radiation hard IP database.
Develop and maintain IP library for future flight system applications.
Establish IP working group
Establish database of available fabrication and design facilities: roadmaps, costs, tool interfaces, availability.
Identify and characterize space rated technologies.
Design, fabricate, test, and evaluate diagnostic structures and circuits used in characterizing foundry process and addressing specific issues of targeted subsystem designs.

X2000 FUTURE DELIVERIES (2^{nd}, 3^{rd}... deliveries)

Although currently planed for 1^{st}, 2^{nd} and 3^{rd} deliveries, X2000 is for the continuing evolvement of technology. It is NASA's intention to establish this program as an ongoing technology development program with deliveries to flight on three-year centers. This paper concentrates on the evolution from 1^{st} through the 3^{rd} deliveries.

During FY '98, a team, lead by C. Salvo, was formed to evaluate technologies, missions and science thrusts appropriate to the X2000 2^{nd} delivery. It was determined that the 2^{nd} delivery should concentrate on small spacecraft platforms and in-situ missions. It was further decided that the goal would be the delivery of a complete proto-flight spacecraft by end of FY '03. Additional goals for the generic spacecraft study were to:

1. leverage off the X2000 1^{st} delivery work,
2. enable early spin out of technology from the System On A Chip (SOAC) program
3. create a bridge between X2000 1^{st} and 3^{rd} deliveries (SOAC based)

Thus, the approach taken was to highly leverage the first delivery investment in:

Micro-controller Soft IP Cores
System Buss/Interface Soft IP Cores
System Software
Power Control and Distribution ASIC design
ASIC design capabilities, technology and infrastructure

The approach also leverages SOAC technologies and JPL technologies such as:

JPL radio on a chip development
SOAC Embedded/Integrated Passives
SOAC Advanced Packaging
SOAC MEMS Micro-Gyro/Accelerometer
SOAC APS for Micro Star Camera and Science Imager
SOAC Complex multi-IP integration technology and infrastructure

In essence, the approach is to:

1. Re-cast the X2000 1st delivery micro-controller, including processor, memory, analog I/O, digital I/O, power system elements, and bus interfaces into fewer, state of the art, ASICs in a commercially available technology. This provides an effective 2 to 3 generation advance in technology allowing a multi-chip set to be implemented as a single ASIC, at significantly reduced power and commensurately improved performance. This aspect of the approach also allows reuse of the X2000 1st delivery system software and software development infrastructure as well as reuse of the hardware/system test and integration facilities/equipment.

2. Implement early versions of some of the X2000 3rd Delivery/SOAC technologies, thus forcing early productization of these technologies and improving probability of success of the, SOAC based, 3rd delivery project.

3. Create standard modules that can be mixed and matched as necessary for a variety of missions.

4. Package the electronics into an extremely compact, rugged and modular form suitable for micro-missions.

5. Integrate units, which traditionally have been kept separate, to take advantage of synergy and modularity, e.g.: create a multi-function aperture serving both science imaging and star-tracking by mounting a beam splitter an separate star-tracker APS in the science imager(s); create a common instrument processor for all science instruments; create a unified ACS processor for both star tracker and inertial navigation; create a unified communication processor for both RF and optical com.

6. Take advantage of the extremely small volume to reduce drive power, provide mass margin for shielding, etc.

In summary, X2000 can be described as a evolutionary process of flight hardware. Figure 3 depicts the how the various component merge to produce the deliverables.

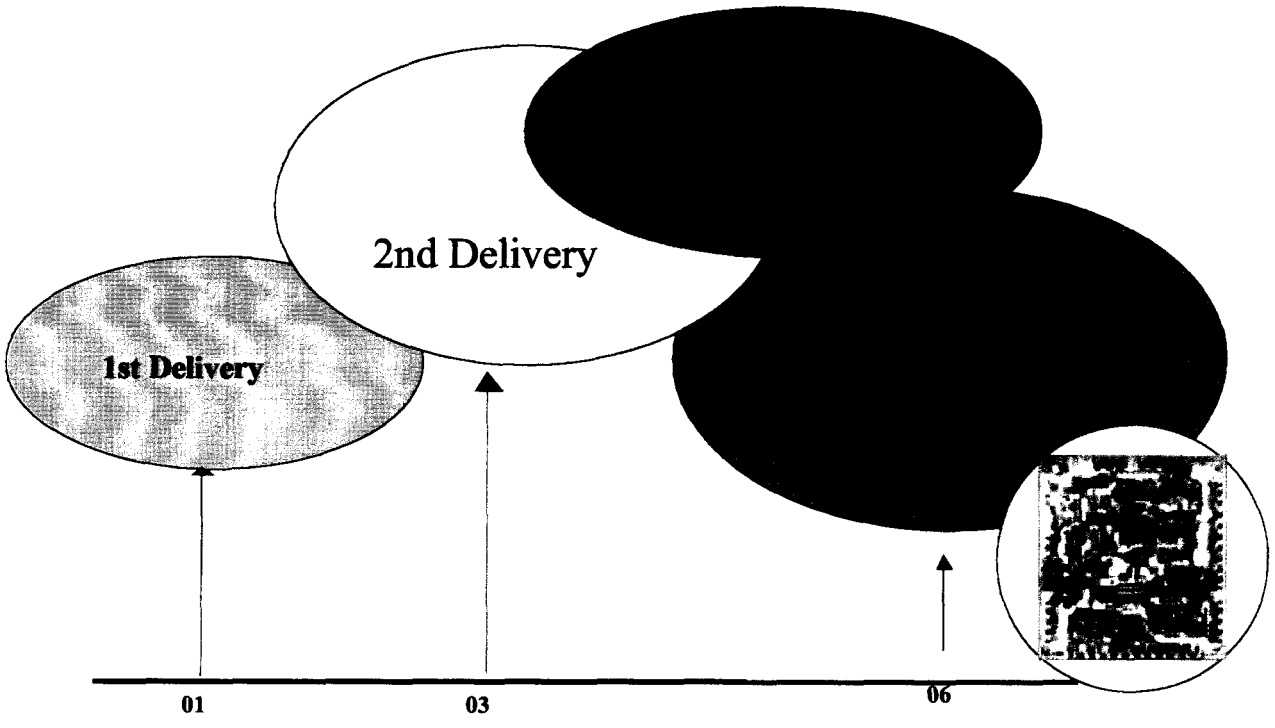

Figure 3. X2000 EVOLUTION

X2000 2nd DELIVERY

As previously mentioned, the 2nd delivery leverages the work done by the 1st delivery and takes advantage of commercially available, but cutting edge, technology and the infusion of SOAC technologies. 2nd delivery is on an evolutionary path starting with 1st delivery and leading to system on a chip. The following discussion describes some of the evolutionary path from 1st to 2nd deliveries.

After reviewing the 1st delivery architecture (Figure 1), we have decided to maintain the same architecture. Which is, to maintain the principles of a distributed and scalable architecture. A number of system attributes were targeted for an evolutionary change. The major changes are outlined below. These changes, targeted for an Engineering Model (EM) delivery by the year 2003, are made possible by deeper submicron processes and other technologies, which will become available.

1. Change the Power Switch Slice from the 1st delivery "centralized" architecture to a distributed (and still scalable) one. This will lead to a "bus ready" approach for spacecraft sensors and instruments.

2. Move the now external PCI bus to be an "internal to the slice" bus. That is - The 1st delivery utilizes PCI connectivity between slices. e.g the Space Flight Computer communicates with NVM and the System I/O (SIO) via the PCI. This is driven by the slice partitioning and lead to a Compact PCI packaging approach. The 2nd delivery seeks to eliminate the need for this "back plane".

3. Repartition the slices to make the above possible.

BUS READY SENSORS

The 1st delivery PSS is a centralized approach which will utilize 16 to 24 I^2C controlled switches on an individual PCI card. Considering the typical spacecraft loads (sensors, valves, pyros etc.) the number, volume and mass of the wires connecting the PSS switches to the loads could become overwhelming. This would become a major concern.

The 2nd delivery approach is to leverage the 1st delivery PSS elements (circuits, ASICs, technology). With a small additional effort the PSS can be reengineered to provide a standardized sensor interface. This interface will enable an I^2C controlled bus voltage ON/OFF command to the lead. It will also provide the I^2C sensor data interface. The intention is to require all sensors to be "bus ready". This will enable a distributed sensor system and will save extensive mass, volume, and integration and test effort. X2000 will provide the sensor interface module to all sensors and instrument that need it. This repartitioning is made possible, and attractive, by utilizing 0.35micron processes which (compared to the 0.8 micron 1st delivery process) enable a gain in functional density and make a single-chip based interface possible.

OTHER SLICES

The 2nd delivery concept is to require all the slices to be bus ready. This will eliminate the requirement for the additional pin-intensive PCI connectivity. Figure 4 below depicts an example for the proposed change. In this example the 1st delivery SFC, which requires both NVM and SIO slices to be functional is repartitioned into two bus ready 2nd delivery nodes. By utilizing the 1st delivery circuit design, IP and ASICs, and with a reasonable additional engineering, the PCI connectivity is eliminated and both the flight computer and the NVM will become bus ready nodes. This repartitioning is made possible by taking advantage of 0.35-micron SOI process and other technologies, which will be available in the appropriate time frame.

Figure 4 is an example for the migration from 1st to 2nd deliveries. As a result, we are eliminating the need for the PCI external "Z direction" connection and consequently the PCI card and cage packaging. In addition, we now reduced the number of slices (in this example) from three to two and each resulting "node" will be functionally more powerful, will consume less power, and will reduce mass and power.

Figure 5 depicts the resulting 2nd delivery Flight Computer Node design.

A similar process is being applied to the rest of the 1st delivery slices. The NVM (None Volatile Memory), the MCS (Micro Controller Slice), the power slices, etc. This process will result in higher functional density, will reduce mass, power and volume, and will reduce the cost of integration and test.

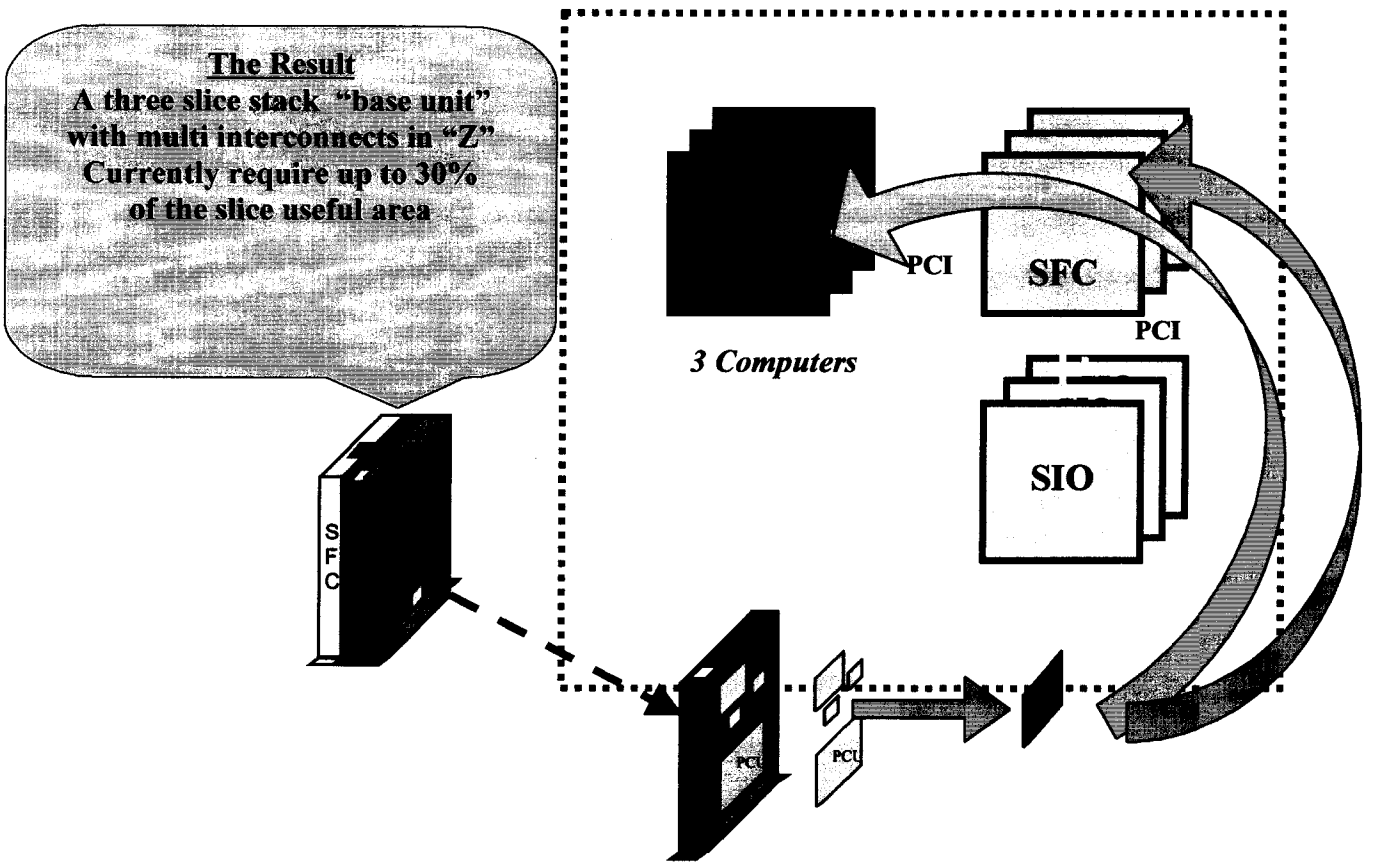

Figure 4. 1st DELIVERY TO 2nd DELIVERY MIGRATION

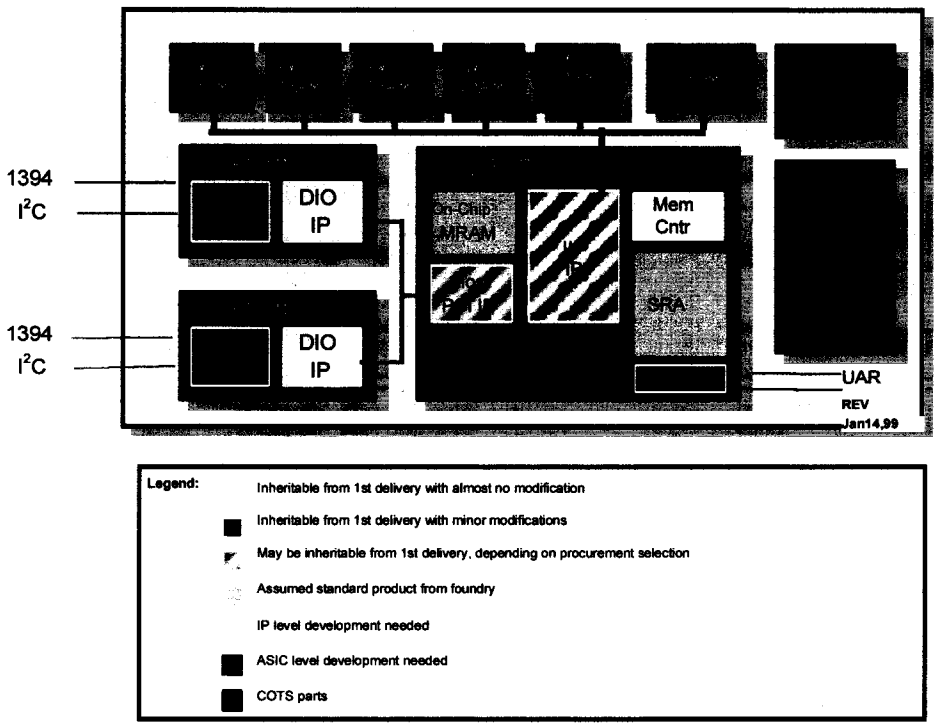

Figure 5. 2nd DELIVERY FLIGHT COMPUTER NODE (FCN)

A similar process is being applied to the rest of the 1st delivery slices. The NVM (None Volatile Memory), the MCS (Micro Controller Slice), the power slices, etc. This process will result in higher functional density, will reduce mass, power and volume, and will reduce the cost of integration and test.

The current 2nd delivery architecture is depicted in Figure 6.

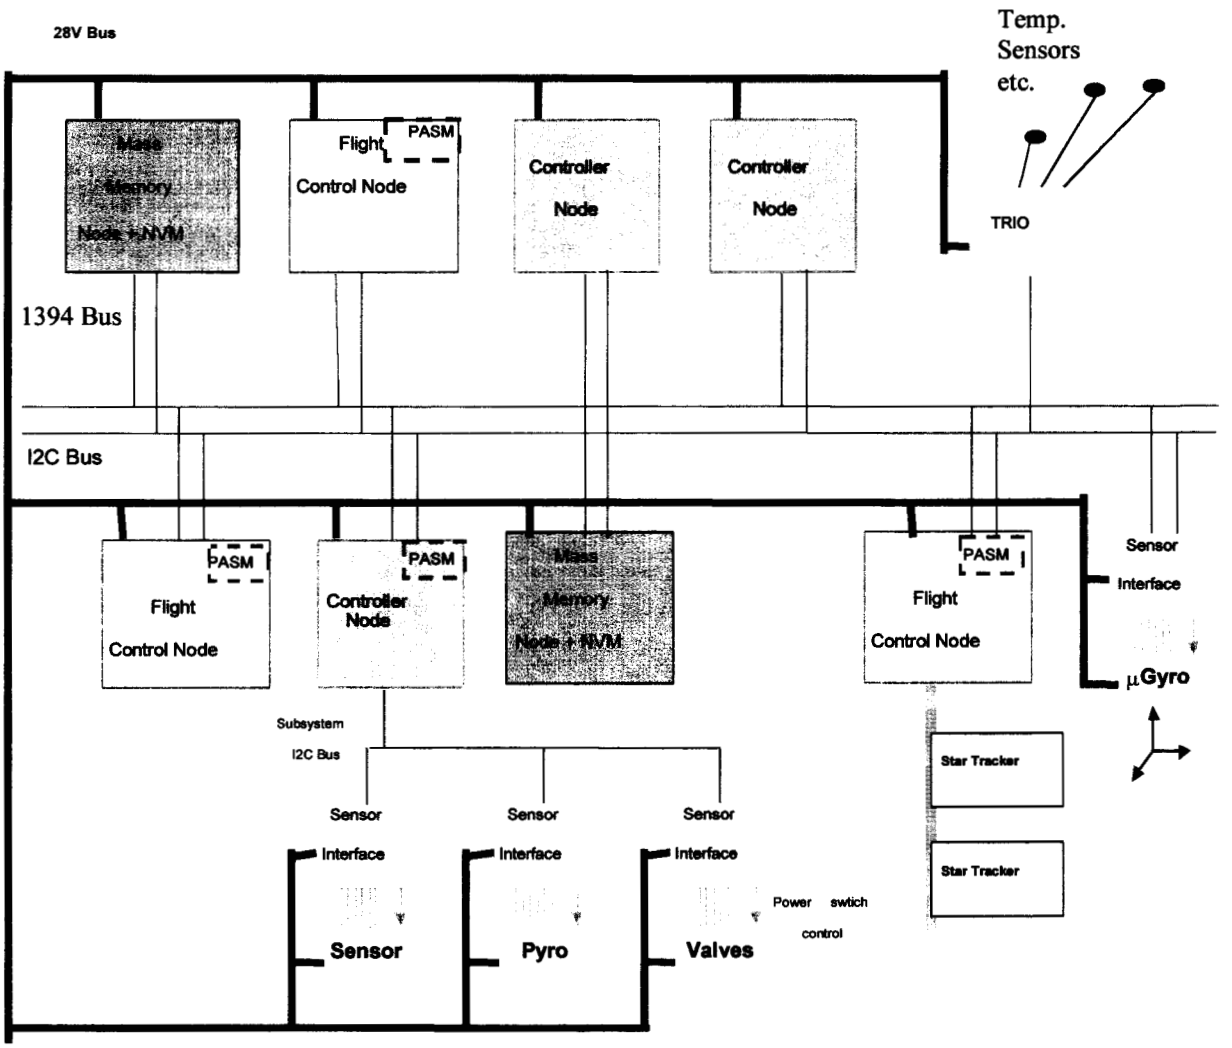

Figure 6. 2nd DELIVERY ARCHITECTURE

X2000 3rd DELIVERY

The objective of the 3rd delivery is to fly a System on A chip. The details and definition are yet to be defined as part of the X2000 plan for FY 99. The undepinning principle is that we must take advantage of higher level integration. This will push the "integration and test" bounderies further and further out and therby save on the high cost and time consumenig I&T tasks. The full implementation of a "System" on a Chip requires the integration of avionics, power management, local power source(s) and power storage, sensors and, in certain cases, actuators. X2000, through the SOAC program is addressing the technologies that will make such possible. A viable "spacecraft on a chip" will one day free us to concentrate on the true goal of space exploration i.e. the science apertures and the collection of scientific data.

The figure below depicts the roadmap for spacecraft evolution through the 2010 timeframe.

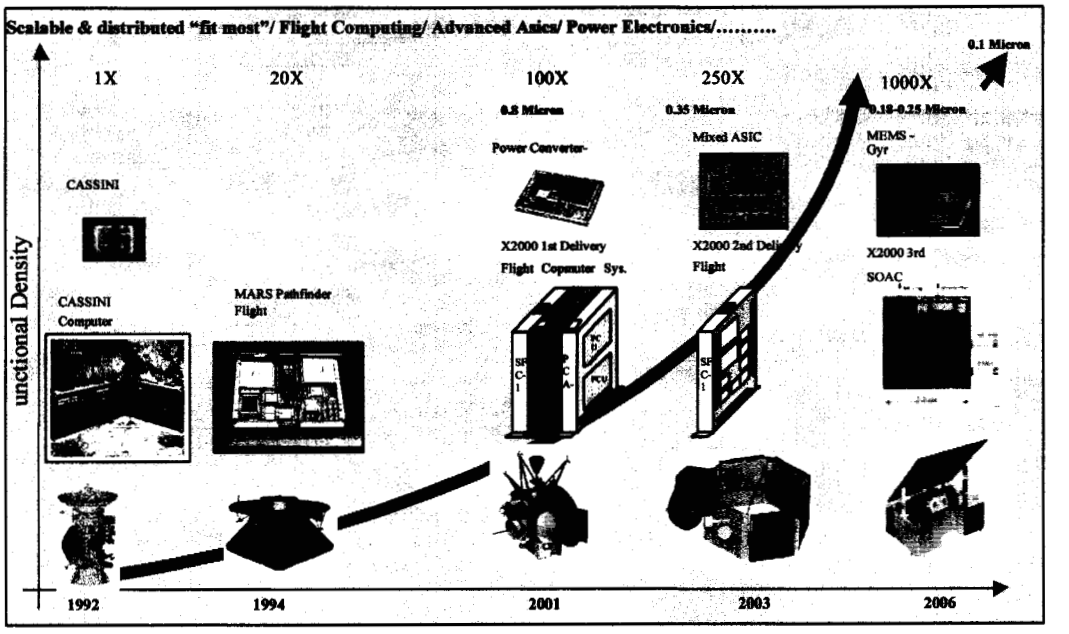

Figure 7. AVIONICS ROADMAP

Figure 8 below is a good summary of the X2000 view.

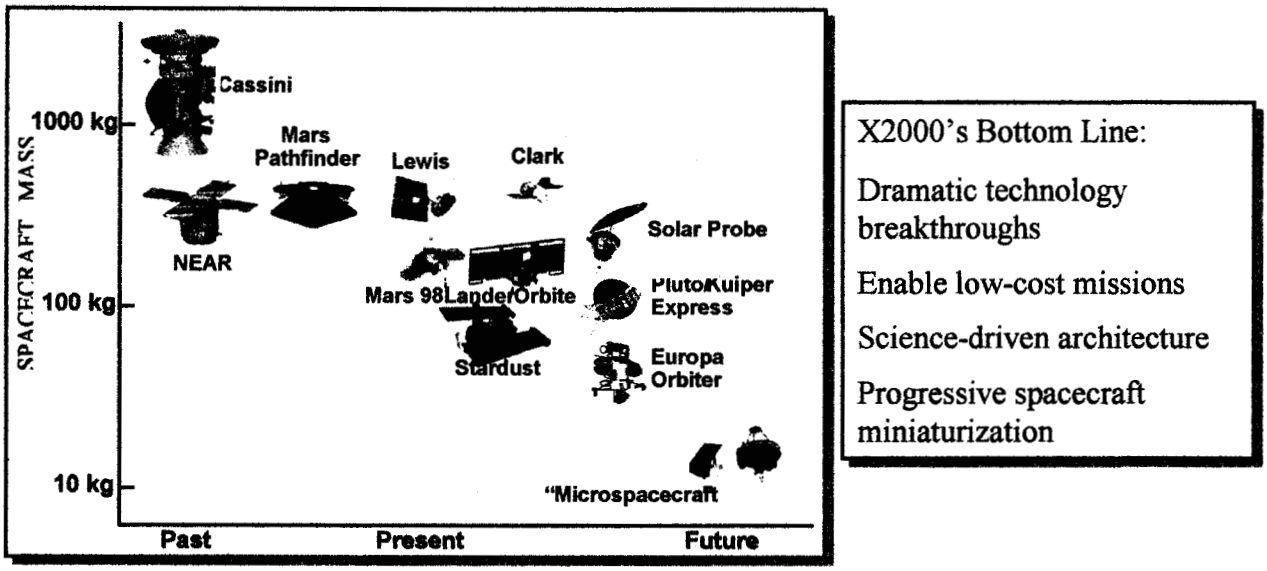

Figure 8. X2000 summary view

Acknowledgments:

Our thanks to the X2000 program, the Center for Integrated Space Microsystems (CISM), and the System On A Chip (SOAC) program for their support and help.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.